

Verifiable Round-Robin Scheme for Smart Homes*

Nisha Panwar, Shantanu Sharma, Guoxi Wang, Sharad Mehrotra, and Nalini Venkatasubramanian
University of California, Irvine, USA.

ABSTRACT

Advances in sensing, networking, and actuation technologies have resulted in the IoT wave that is expected to revolutionize all aspects of modern society. This paper focuses on the new challenges of privacy that arise in IoT in the context of smart homes. Specifically, the paper focuses on preventing the user's privacy via inferences through channel and in-home device activities. We propose a method for securely scheduling the devices while decoupling the device and channels activities. The proposed solution avoids any attacks that may reveal the coordinated schedule of the devices, and hence, also, assures that inferences that may compromise individual's privacy are not leaked due to device and channel level activities. Our experiments also validate the proposed approach, and consequently, an adversary cannot infer device and channel activities by just observing the network traffic.

CCS CONCEPTS

• **Security and privacy** → **Security protocols**; *Mobile and wireless security*; *Domain-specific security and privacy architectures*; *Social aspects of security and privacy*;

KEYWORDS

Internet of Things; smart homes; user privacy; channel and device activity; inference attacks.

ACM Reference Format:

Nisha Panwar, Shantanu Sharma, Guoxi Wang, Sharad Mehrotra, and Nalini Venkatasubramanian. 2019. Verifiable Round-Robin Scheme for Smart Homes. In *Ninth ACM Conference on Data and Application Security and Privacy (CODASPY '19)*, March 25–27, 2019, Richardson, TX, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3292006.3300043>

1 INTRODUCTION

The IoT devices are quickly becoming a pervasive and integral part of modern smart homes [9]. The homeowner, typically, possesses a heterogeneous set of devices ranging from wearable devices, information/entertainment devices to smart home appliances. These

*This material is based on research sponsored by DARPA under agreement number FA8750-16-2-0021. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government. This work is partially supported by NSF grants 1527536 and 1545071.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CODASPY '19, March 25–27, 2019, Richardson, TX, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6099-9/19/03...\$15.00
<https://doi.org/10.1145/3292006.3300043>

devices provide comfort/assisted-living and/or improve sustainability, reduce costs, and reduce carbon footprint. For example, a Belkin Wemo switch can automatically switch lights on/off and open/close window shades based on the sunlight and time of the day. Likewise, dampers in the AC vent can be partially/fully opened/closed to modulate airflow. Other devices popular in smart homes include Nest cameras, smart door locks, Lenovo Smart Assistant, Amazon Echo, Echo dot, Echo show, Alexa, Philips-Hue Bloom/Lightstrip Plus, SteriGrip self-cleaning door handles, Unico smartbrush, Sensus Metering Systems, and Logitech Circle 2 among others.

While the emerging smart home devices provide significant benefits, the support for security in such devices is often limited to the security offered by the original equipment manufacturer (OEM). Lack of strong end-to-end architecture for security has led to devices being vulnerable to a variety of attacks. For example, McAfee Labs [1] found that the well-known Wi-Fi-enabled Wemo Insight Smart Plug has critical security vulnerability due to Universal Plug and Play (UPnP) protocol library it uses, which, due to design flaws, enable attackers to execute remote codes on this smart plug. Note that this attack is not just limited to disturbing smart plug's normal operations such as shutting it down unexpectedly, but could also use the smart plug as an entry point for a larger attack in the network. Further, [1] showed the usage of a compromised WeMo switch as a middleman to launch attacks against a TCL smart TV.

The privacy vulnerabilities introduced by smart home devices are even more challenging. IoT devices capture, store, share, and (depending upon the underlying computational architecture) out-source personal data that can lead to inferences about individual's habits, behavior, family dynamics etc. Challenges arise since privacy leakage can occur through direct data leakage, as well as, through inferences based on device actuations, interactions, and schedules. For instance, the timing of the actuation of a coffee machine, if leaked can allow an adversary to determine when a family wakes up. Likewise, locking and unlocking schedule of door locks can enable leakage of the time when no one is at home, etc. While privacy challenges from data leakage can be prevented by encrypting device data and network payloads, inferences about device actuation and schedule are significantly more complex to hide due to leakage from network traffic patterns at the channel level, at the hub/router level or at the cloud level.

For example, Figure 1 shows the channel traffic generated by three different devices. The figure clearly shows that each device generates a very distinct traffic pattern and, the adversary, having access to the channel traffic can figure out which device is activated leading to potential inferences about user's personal habits. Note that such an inference, since it is independent of the actual network payload, is not prevented by encryption.

Inferences from monitoring channel traffic can also arise due to the characteristics of the current network protocols. For instance, in the widely used 802.11 Wi-Fi protocol, while a message payload is encrypted in a password-protected network, the MAC addresses

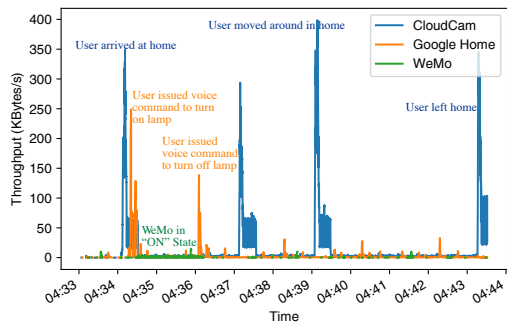


Figure 1: Channel activity for home devices.

The figure above shows the channel activity for three home devices: CloudCam, Google Home, and Belkin WeMo. The CloudCam shows a peak in the channel activity (up to 400 KB/s traffic rate) as the user enters the home, moves inside the home or exits the home. The Google Home shows a peak in the channel activity (up to 250 KB/s traffic rate) whenever a user initiated a voice command for the light bulbs to turn on/off. Similarly, a bi-state WeMo switch peaks during the on state and creates a channel activity lesser than 20 KB/s.

of both the sender and receiver are in cleartext. This is to prevent requiring every potential device on the network to have to decrypt a message just to determine if the message is intended for the device. The leakage of the sender/receiver MAC address, coupled with the fact that manufacturer’s information is commonly encoded in plain-text device identifier, can lead to leakage of the identity of the device from the network traffic, which, in turn, can lead to an attack on user’s privacy. For example, the MAC address of Amazon CloudCam security camera used in our experiment to transmit video footage over Wi-Fi is F0-81-73-23-CC-75. The first 3 bytes of such a MAC address (e.g., F0-81-73) can be searched in the publicly accessible IEEE Organizationally Unique Identifier (OUI) [2] dataset to find the vendor related information (e.g., Amazon device). Furthermore, by monitoring the device’s traffic patterns and the fact that Amazon only manufactures a limited range of devices (e.g., Kindle, CloudCam, Echo, etc.), it is easier to infer the device type by merely overhearing the traffic.

In this paper, we study privacy leakage that may occur from device activity and network channel traffic analysis and develop protocols that can be used to prevent such leakages. We focus, in particular, on device workflows that are common in smart homes. By a device workflow, we refer to a coordinated sequence of device actuations. Device workflows may arise in a *triggered* (or *synchronized*) manner or in a *scheduled* manner. The synchronized workflows arise as a result of one device resulting in an actuation of the other. For instance, sensors determining occupancy change in a part of the building may result in HVAC controls /AC vents to be redirected to the occupied areas and to close other vents that cover areas with no occupancy. Likewise, light intensity sensors coupled with thermal sensors may detect the amount of sunlight entering the room and accordingly lower/raise the sunshades based on the homeowner’s preference. A scheduled workflow, on the other hand, is scheduled actuation of a set of devices that occur at specific time intervals of each other based on a schedule. For instance, switching on a coffee machine at a specific time in the morning followed by warming of the car seats a given time interval following that, and then opening/closing of door locks following the actuation of the car seat warmer might be on a schedule. A more elaborate example

of a scheduled workflow could be a homeowner’s routine related to returning. A homeowner may schedule the smart car to self-drive to home to initiate the workflow. Fifteen minutes after the start of the workflow, the heating/cooling system may start off to ensure that the home is at the comfortable temperature on arrival. Likewise, half an hour after the trip starts the oven may be set to a pre-heat and the laundry machine turns on if the load is detected.

We focus on the scheduled workflows in this paper since such workflows require hiding the identity of the devices being actuated but also their schedule. As will become clear, mechanisms to prevent leakage for scheduled workflows will also prevent leakage from triggered workflows. Furthermore, the scheduled activities can lead not just to adversary learning user’s past behavior but also their future activities which can lead to more severe consequences.

The problem. This paper deals with a problem of avoiding inference attacks on the scheduled workflows in a home network. The workflows can be identified through *coupling* between the *channel* and *device activity*. Basically, there are two crucial concepts that are subject to privacy violations: workflow (i.e., the specific order of device actuation) and workflow execution (i.e., duration in which the devices coordinate, and the resulting device actuations unfold). Our problem statement considers hiding both the workflow and the execution of the workflow.

The privacy violations can occur as a result of two threats: first, overhearing the channel activity as a means to infer device activity pattern, second, accessing the device temporarily and be able to analyze the state of workflow execution. In the latter case, the adversary can read the sent/received messages or the internal state of the device. Both of the above threats may assist the adversary to predict users’ activities, such as presence/absence, arrival/departure, and localization etc.

Contributions. Our contributions are twofold:

- (1) A new architecture for in-home communication among the devices and the hub through passing a token carrying commands. The token passing communication model decouples the channel and device activities so that the devices interact with the hub and the other devices without revealing the communication pattern. In addition, this architecture is also useful for secure data upload from devices to the hub, while also hiding the device footprints that has generated the data.
- (2) We provide an owner-defined pre-scheduling mechanism for all devices that are connected with the hub in a pre-defined topology. The proposed approach uses a single message transmission for all N devices while ensuring that the in-home communication remains peakless. The scheduling mechanism is secure against a computationally unbounded adversary and, also, verifies the delay between each device actuation.

Outline. The paper proceeds as follows: Section 2 provides the system setting, the adversarial model, security goals, and design requirements. Section 3 provides our proposed scheduling algorithm for home networks. Section 4 provides proofs of security and privacy. Finally, Section 5 provides an experimental evaluation of the approach. All notations are given in Table 1.

| Notations | Meaning | Notations | Meaning |
|-------------------------|--------------------------------------|----------------|----------------------------------|
| O | A homeowner | H | Hub |
| O_{id} | Owner's identity | H_{id} | Hub identity |
| D | Device | D_{id} | Device identity |
| O_{PK} | Owner public key | O_{SK} | Owner secret key |
| H_{PK} | hub public key | H_{SK} | hub secret key |
| D_{PK} | device public key | D_{SK} | device secret key |
| c_l | partially ordered l commands | m^i | puzzle message for i th device |
| n | modulus | a | random chosen integer |
| \hat{i} | time complexity of puzzle | t' | time to decrypt command |
| S | capacity of puzzle solver | S' | enhanced capacity |
| $Sign(O_{SK})$ | Signature using secret-key O_{SK} | N | number of devices |
| E_k | encrypted key k | E_z | encrypted command z |
| k_s | static key | $n P_n$ | n permutations |
| $\phi(n)$ | Euler's totient on n | b_o | overwritten data bits |
| b_g | device generated data bits | b_r | random data bits |
| p | large prime number | q | second large prime number |
| t_{val} | command validity time | t_{cur} | current time |
| t_{rcv}^i | token receiving time | t_{fwd}^i | token forwarding time |
| t_{beg}^H | token round beginning time at hub | t_{end}^H | token round ending time at hub |
| t_{diff} | allowed clock drift time | t_{com} | total computation time |
| $t_{com}^{\mathcal{A}}$ | puzzle computation time by adversary | ϵ | negligibly small value |
| \mathcal{A} | adversary | \hat{a} | malicious commands |
| \mathcal{T} | token | \mathcal{H} | One-way hash function |
| <i>Data field</i> | data upload field | b_{toggle} | toggle bit string field |
| \mathcal{R} | partial order | \mathcal{E}' | partially ordered set |

Table 1: Notations

2 PRELIMINARIES

This section presents the system model, the adversarial model, inference attacks on the user privacy, an overview of our proposed approach to prevent inference attacks, design requirements, and building blocks of the proposed algorithms.

2.1 The Model

Network assumptions. We consider a homeowner, O , who owns a collection of N (D^1, D^2, \dots, D^N) heterogenous smart home devices that provide different functionalities to the owner. Each device D^i has a unique identity, denoted by D_{id}^i . All devices might possess heterogeneous hardware/software underneath, and be located on different spatial (devices that are not in the line-of-sight) dimensions. These ad-hoc devices can shift in space in the smart home, and hence, might have a different set of peer devices at different time intervals. We assume that each device possesses a read-only hardware clock, and due to the ad-hoc nature of devices, we assume a clock drift within the bound t_{diff} such that two clocks cannot differ beyond the t_{diff} amount of time.

The owner initializes the devices and a controlling hub, H , using proper security mechanisms. We will list our assumptions about the underlying security mechanism below. In our model, the network is configured as a ring topology, which poses an ordering among devices, unlike the model that the current smart home devices use, where the owner communicates directly to the desired device via the hub. This ring topology could be built directly among devices and hub, if the communication protocols they use have P2P communication capability, like Zig-Bee, BLE, and Wi-Fi. Alternatively, it

could be built as an overlay on top of a star topology network, like Wi-Fi infrastructure mode. In this case, if a device tries to forward the message to the next device in the proposed ring topology, then it needs to first send the message to the hub, and the hub, then, directly forwards the message to its next device. Note that we do not discuss a failure-resilient ring topology and existing fault-tolerant schemes can be leveraged here.

The owner sends workflows to the home devices through the hub. After receiving a workflow, the device gets actuated, stores its corresponding command, and forwards the workflow received from the previous device to the next device in the topology. After executing the command, the device may generate the data (for example, Nest camera starts recording and sends data whenever motion is detected). We use the ring topology to send this data to the hub that may be stored at the hub or may be transmitted to the cloud.¹ In this paper, we use the words ‘command execution’ and ‘workflow execution’ interchangeably.

Token. In our ring topology, we circulate a token that has three fields: (i) command field, which carries a computational puzzle and the workflow, (ii) data field, which carries the data generated by devices to deliver to the hub, and (iii) toggle bit string, which is used to indicate which device has generated the data to the hub. Details of the token are given in Section 3.

Security assumptions. Each device, hub, and homeowner possess the corresponding signing key-pair, i.e., (D_{SK}^i, D_{PK}^i), (H_{SK}, H_{PK}), and (O_{SK}, O_{PK}), respectively. We do not assume an arbitrary behavior from the owner or the hub. The *homeowner and hub* mutually verify the identities of each other through digital signatures in order to build the trust between them. Therefore, the hub and the owner trust each other, and an adversary cannot compromise either the homeowner or the hub. Further, *the hub and devices* build their own trust that is also based on the knowledge of a certified public-private key-pair of home devices.

2.2 Adversarial Model

The devices execute user-defined commands or workflows, as mentioned previously. The adversary wishes to learn the commands or workflows and the execution of workflows based on the encrypted network traffic to infer the user privacy. This type of adversary is similar to the adversary considered in [6, 14]. Thus, the adversary has access to the secure (encrypted) messages flowing among the devices and the hub/homeowner. Further, we assume that the adversary knows the number of smart devices in the home. Based on this information, the adversary aims to learn: (i) the device activities, and (ii) coupling between the channel and the device activities. However, the adversary cannot inject any fabricated messages over the channel to assess the state of the devices.

Further, we assume that the adversary may gain a short-term physical access² to the device, and hence, can retrieve the device

¹Recall that we are not dealing with how the data will be transmitted from the hub to the cloud without revealing anything. Our solution hides any activity within the home, i.e., how the data will be transmitted by the device to the hub without any privacy violations.

²For example, an inspection authority has to visit the home for a periodic inspection in the absence of the homeowner. In this case, the inspection might be related to any leakage detection, maintenance issues, insurance issues, etc. This short visit to the home for an inspection allows them to monitor and check home devices as well. However, based on our proposed solution, those inspecting authorities would not be

state or messages. The objective behind gaining a short-term access to any device is to predict the future workflows of devices to infer the user activity. Our objective is to prevent the adversary to know (i) which device has received the messages at which time, (ii) when would a device execute the command, and (iii) which devices have executed the command at which time³.

Adversarial view and inference attacks. When the user wishes to execute any command at a smart home device, an adversary knows which device received the message from the user at what time due to the network traffic generated by the user. Note that this information is revealed, because in network protocols such as 802.11 MAC addresses or the device identifiers are transmitted in cleartext, and only the payload is encrypted, as mentioned in Section 1. Further, the device may also produce some data in response to the requested message, and this also reveals to the adversary which device has generated the data at what time. We refer such information as the adversarial view, denoted by AV : $AV = In_c \cup Op_d$, where In_c refers to the command given to the device at some time and Op_d refers to the data generated by the device after executing the command.

| Users command | Adversarial view | |
|---------------|--------------------|--------------------|
| | In_c | Op_d |
| For D^1 | $E(c_1), D^1, t_1$ | No |
| For D^2 | $E(c_2), D^2, t_2$ | $E(d_2), D^2, t_3$ |

Table 2: Adversarial view.

For example, consider that there are two devices, say D^1 and D^2 , in the home. In Table 2, the first row shows that the user transmits a command to the device D^1 . Though the device D^1 receives this encrypted command, denoted by $E(c_1)$, the adversary knows that a command is received at time t_1 by the device D^1 and the device D^1 has not generated any data in response the command. The second row shows that the adversary knows the device D^2 receives an encrypted command $E(c_2)$ at time t_2 and generated encrypted data $E(d_2)$ at time t_3 . Hence, simply based on the above characteristic of the arrival of a message and generation of data, the adversary can determine which device was actuated.

2.3 Preventing Inference Attacks: overview of our Approach

In order to prevent inference attacks, we develop an approach that decouples the device and channel activities. In short, the approach provides an ability to pre-schedule a set of commands for home devices, where the homeowner defines: (a) what should be the workflow/schedule of home devices, and (b) when should the devices execute a workflow. Informally, the proposed approach works as follows:

- (1) The owner invokes the hub by sending an encrypted schedule or workflow of the devices. Here, the hub authenticates the owner to validate the encrypted schedule. This step provides a guarantee that no channel spoofing or message re-transmission have occurred.

able to analyze the current state of the devices or the activity pattern of the devices in near future.

³The command execution could enable the device to produce visible or auditory cues such as blinking lights or machine being activated which, in turn, may leak the state of the device. Such inferences from physical cues are outside the scope of the paper. We assume that the adversary does not have access to such device data.

- (2) After a successful authentication phase, the hub creates a token (\mathcal{T}) to circulate the schedule to be executed by devices. This token rotates continuously in the topology. Note that in the token, device identifiers or MAC addresses are also encrypted. Further note that whenever the user wishes to transmit a schedule to devices, the immediate next round of token originated by the hub carries the encrypted schedule, and after that, the encrypted schedule is replaced by a random message to maintain the token size constant.
- (3) On receiving the token, each device retrieves the encrypted schedule that carries device-specific commands. Each device must complete a computation task before executing the original command. This computation task is referred to as a puzzle throughout the paper. Then, the device must check the puzzle validity time t_{val} , by using the current time t_{cur} , and the allowed clock difference t_{diff} . If the timer has not expired yet, the device decrypts the message, executes the puzzle to retrieve the real command to be executed.
- (4) As soon as a device finishes the command execution, it may generate data to be uploaded at the hub (as mentioned in Section 2.1). Now to hide which device has generated the data, each device follows a request-based approach, where a device D^i flips the i^{th} bit inside one of the fields of the token to indicate the need to upload freshly generated data in an anonymous manner. As a result of i^{th} bit flipping, the hub knows that the device D^i has requested the data upload, and in the next token cycle, the device D^i appends the data inside a dedicated field of the token.

Note that since a constant size token flows regularly in the topology of the home devices, an adversary observing the network traffic cannot distinguish which device has received a command at which time (due to step 2) and which devices have generated the data (due to step 4). Hence, based on this approach, the adversarial view for each round of token has the same information, which prevents inference attacks based on the devices and channel activities.

2.4 Security Goals

This section describes the security properties for preventing any inferences about the workflow and their executions from the adversary. Let us assume that an adversary knows some auxiliary information about the devices and the topology such as the number of devices and the types of devices. However, this auxiliary information does not increase the probabilistic advantage that an adversary gains over any instance of the protocol. In particular, an adversary cannot reveal the workflow or the execution time of the workflow, i.e., which devices execute the command or when does a device execute the command. The probabilistic advantage of an adversary, denoted by $Adv(\mathcal{A})$, is derived through the security properties given below.

Authentication is required during the workflow release from the homeowner to the hub. This would require a mutual authentication between the (mobile device held by the) homeowner (to dispatch the workflow) and the hub (to circulate the workflow anonymously). Note that establishing the shared secret between the homeowner and the hub is a one-time process, which is carried each time the homeowner invokes a new workflow. Here, the homeowner produces a signature, say $Sign$, on the ordered commands (c_1) in any workflow by using its secret-key O_{SK} . Thus, the hub must reject any other messages signed by a different key, say $O_{SK'}$.

$$\Pr[(O_{SK}, c_l) \rightarrow \text{Sign}] \geq 1 - \epsilon$$

Note that ϵ is negligibly small and an adversary cannot produce a verifiable signature Sign on c_l by using $O_{SK'}$ instead of O_{SK} .

Anonymity is required during the consistent circulation of the encrypted commands such that (i) no channel activity can be mapped to a device activity, and (ii) no inference on device activity can be mapped to the device generated data, i.e., which device is sending data at a specific time. As shown below, the probability of distinguishing two different tokens ($\mathcal{T}, \mathcal{T}'$) each carrying different messages (m^i, m^j) for different devices (i, j) is negligibly.

$$\Pr[\mathcal{T}(m^i)] - \Pr[\mathcal{T}'(m^j)] < \epsilon$$

Similarly, the probability of distinguishing a token \mathcal{T} carrying the random data b_r or carrying the overwritten data b_o , is negligibly small. Therefore, a token carrying the random data inside the data field and another token carrying the overwritten data (after data generation) inside the data field are indistinguishable, hence, solely based on the token data field no inferences can be derived.

$$\Pr[\mathcal{T}(b_r)] - \Pr[\mathcal{T}(b_o)] < \epsilon$$

Verifiable delay An adversary cannot infer the information about the device execution ahead of time. Let \hat{t}_i be the time a device would execute even when the adversary has temporary access to the device.

$$\Pr[t_{com}^{\mathcal{A}} | \text{state}] \approx \Pr[t_{com}^{\mathcal{A}}]$$

The probability of an adversary finishing the computation task earlier, when it gains temporary access to the device state, is approximately same as when the adversary does not have access to the device state. In addition, an adversary cannot outpace a device that requires \hat{t}_i time to complete the computation task, i.e., $\Pr[t_{com}^{\mathcal{A}} < \hat{t}_i] < \epsilon$.

An adversary gaining access to the device cannot retrieve the information required for device actuation, i.e., the computational task to be executed prior to its actuation. We have described a game-theoretic approach in Section 4 that shows the overall probabilistic advantage of an adversary is negligibly small.

2.5 Challenges and Solutions

Implementing reliable ordering for device actuation is to provide a secure and self-executing state of devices at a pre-defined time for protecting the owner privacy is deceptively non-trivial. Below, we discuss the challenges we encountered and describe how we addressed them.

C1. Anonymous trigger from the hub to devices. As mentioned before, we need to mask the channel activities, device activities, and the coupling between both the channel as well as the devices. Note that in this context, the encryption techniques merely hide the meaning of the message across the channel, not the fact that to which device this message belongs to, and hence, it reveals the user activity. In addition, the solutions-based on *traffic shaping*⁴, which incur excessive communication and latency overhead, also fail to decouple the device to channel activities.

Solution. To address this challenge, the distribution of user-commands to each device in the home network is based on a pre-defined topology, e.g., ring, where a token rotates continuously

⁴The traffic shaping solution keeps the constant traffic rate based on a threshold, such that any excessive traffic above the threshold is delayed through a buffer and below the threshold requires additional dummy packets.

within the one directional (1-D) ring topology.⁵ Thereby, channel activity remains consistent and independent of the devices actuated as a result of the workflow. In our context, the token (\mathcal{T}) has three fields: (i) command field, which contains the encrypted commands corresponding to each device, (ii) data field, which contains the device generated data, and, (iii) a toggle bit string field, which contains a N bit string, where each bit denotes a unique device in the topology. This toggle bit string is used to indicate that a device is interested in uploading the freshly generated data during upcoming token arrival at the device. Since the order of device actuation reveals crucial information about the user activity inside the home, our token-based solution guarantees a secure ordering among devices while executing the commands. In addition, the device actuation is controlled in a manner that a recipient device itself cannot pre-decode and/or pre-poned the command execution.

C2. Command execution and verifiable ordering. After masking the channel activity through a constantly rotating token, the next step is to have a verifiable ordering of command execution at each device. Each device receives a command through the token. Now, these devices can decrypt and execute the command immediately. However, it again enables the channel to device activity mapping. Therefore, the next challenge is to insert an artificial delay between a device receiving the commands, and then, executing the commands at an appropriate time, without relating to any specific channel activity. The artificial delay enables a correct execution order at each comparable or non-comparable device⁶.

Solution. The protocol message from the homeowner to hub includes the ordered commands (c_l) that pass through a device D^i to another device D^{i+1} using the anonymous token circulation. The token is encrypted using a shared symmetric key between the hub and the devices in the topology. The recipient device D^i retrieves the encrypted command (m^i) from the command field of the token and begins with a puzzle computation. Note that the devices do not execute the commands immediately after receiving (time-locked) commands. In particular, these devices wait for a pre-defined amount of time before executing the command, such that neither the artificial delay at each device can be known by the adversary in advance, nor the devices can control this artificial delay to postpone or preponed the scheduled commands. However, this waiting period is not idle, and the devices resume on some computational task.

C3. Anonymous response from devices to the hub. Once the devices have received the commands in an anonymous manner, they may generate some data as a result of command execution. However, uploading this data immediately would reveal the device activity patterns. Clearly, this periodic channel activity (in the form of traffic) relates to a specific device that had executed the command recently. Therefore, the upstream data upload on the hub should be anonymous too.

⁵In order to leverage a continuous channel-activity as a means to hide the actual channel-activity, the ring topology is efficient as compared to star alignment. In addition, one can also use the mesh-topology with anonymous-routing that requires asymmetric-key cryptography overheads at each relay-node while sending token between a source and a destination.

⁶The comparable-devices are those that are defined under certain relative order in a workflow. The incomparable-devices are independent and are not restricted under any relative order with respect to other devices.

Solution. In our scheme, each device is required to send a request for data upload to the hub using the rotating token, which contains a toggle bit string field of size N bits, where each bit represents a unique device in the topology. The i^{th} toggled bit indicates that the device D^i has generated the data and is ready to anonymously transmit the data by using the data field of the token. The data field is used to carry the device generated data without revealing the data and the sender of the data. The data field contains random data as long as there is no request from the devices to send the data to hub. The details of data upload phase from device to hub will be clear in step 4, given in Section 3.

2.6 Building Blocks

This section provides a brief overview of basic building blocks used to develop the proposed solution as detailed in Section 3.

RSA puzzles. The verifiable delay regarding device actuation is based on cryptographic RSA puzzles [12]. These time-bounded puzzles are useful for the applications that require security against the hardware parallelization attacks (i.e., bypassing a security solution by running a mathematical problem on different hardware in parallel) through Application Specific Integrated Circuits (ASIC). Accordingly, the puzzle solution is based on inherently sequential operations such as *modular exponentiation*. Let us assume that t' be the time to release an encrypted message. Also, a device is capable of computing S number of square operations modulo n per second. Thus, the puzzle requires sequenced exponentiations of $(a^{2^i} \bmod n)$ where $n = pq$ is publicly known RSA modulus and \hat{t} , p , q and $\phi(n) = (p-1)(q-1)$ remains secret. In particular, $\hat{t} = St'$ denotes the difficulty level of the puzzle for a specific device. Therefore, the computation of this modular exponentiation operation requires either the inherently sequential execution of these operations *or* to solve the integer factorization problem.

In the proposed scheme, the RSA puzzles enable an artificial yet verifiable delay with respect to command execution. To compensate this delay an adversary must know the private key of a device and then invest the same time as the victim device was supposed to invest in, for the puzzle computation. In particular, the adversary can always lengthen the delay (which is easily detectable during the puzzle validity check), but cannot shrink the delay due to inherently sequential operations.

Order-preserving bijection. The order-preserving bijection guarantees an instance of the totally ordered set elements as derived from the partially ordered set of the same elements. In particular, the bijection provides a unique sequence of the totally ordered set elements. Let us assume there is a partial order relation $\mathcal{R} = \{\leq\}$ on set elements $\mathcal{E} = (e^1, e^2, e^3, e^4)$ such that \mathcal{R} is: reflexive, i.e., $e^i \mathcal{R} e^i$; antisymmetric, i.e., if $e^i \mathcal{R} e^j$ and $e^j \mathcal{R} e^i$ then $e^i = e^j$; and transitive, i.e., if $e^i \mathcal{R} e^j$ and $e^j \mathcal{R} e^k$ then $e^i \mathcal{R} e^k$. In addition, the partially ordered set $\mathcal{E}' = ((e^1, e^2), (e^3, e^4))$ under relation $\mathcal{R} = \{\leq\}$ has a unique minimal and maximal element. Therefore, an order-preserving bijection generates a linear ordering of elements in set \mathcal{E}' . Essentially, this linear extension generates a permutation order of the elements in a given partially ordered set \mathcal{E}' . All of those permutation sequences in which e^i appears before e^j given that $(e^i \leq e^j) \in \mathcal{E}'$ are a valid candidate as per any totally ordered set element sequence. The reliable ordering of device actuation

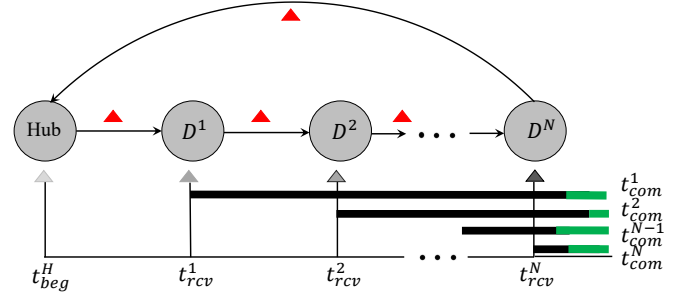


Figure 2: Device ordering on the timeline.

is guaranteed through the linear extension of the owner-defined schedule even when the application host is unavailable.

3 DECOUPLING CHANNEL ACTIVITY FROM DEVICE ACTIVITY

This section provides the details of our proposed protocol for decoupling channel activities from device activities. First, we illustrate the device-to-device interaction through an example as below:

3.1 Example

Consider N number of devices (D^1, D^2, \dots, D^N) that are connected through a hub (H), as shown in Figure 2. Let s_{on}^i (and s_{off}^i) be the on (and off) state of a device i . The owner O can create a partial ordering for devices such as $\langle (D^1, D^2, D^3), (D^4, D^5), \dots, (D^{N-1}, D^N) \rangle$ based on their states, e.g., $\langle (s_{on}^1, s_{off}^2, s_{off}^3), (s_{off}^4, s_{on}^5), \dots, (s_{off}^{N-1}, s_{on}^N) \rangle$ that shows the device D^1 must change its state to on, i.e., s_{on}^1 , before the devices D^2 and D^3 change states to off, i.e., s_{off}^2 and s_{off}^3 . Similarly, the device D^4 must change its state to off, i.e., s_{off}^4 , before the device D^5 changes its state to on, i.e., s_{on}^5 .⁷ After creating the partial order of devices, the owner sends a message to the hub that sends the message (shown in red color) to one of the devices, as shown D^1 in Figure 2. We refer to the message from the hub to devices as a token. Each device i receives the token at time t_{rcv}^i , forwards the token, and begins computation at time t_{com}^i . The bottom part shows when each device receives the token in a sequence as $(D^1, D^2, \dots, D^{N-1}, D^N)$. Note that due to user-defined partial order of device actuation $\langle (D^1, D^2, D^3), (D^4, D^5), \dots, (D^{N-1}, D^N) \rangle$, the devices across the partial orders are mutually incomparable.⁸ In Figure 2, thick black line (for each device) shows that the device is having the token and waiting for the predefined time (given in the token) for its activation, and the green line shows when the devices start working.

3.2 Verifiable Ordering Protocol

The Section formally defines the verifiable ordering protocol and a detailed description for each step as below.

Definition 3.1 (Verifiable Ordering Protocol). The verifiable device ordering protocol is a tuple $(param, puzgen, P_D^i, P_O)$ of four polynomial-time algorithms such that:

⁷For the sake of simplicity this example includes the ordering between the devices and the corresponding states. However, throughout the paper, our focus is to order time intervals for devices' actuation.

⁸In case, the user finds a change in his/her schedule, then another remote command can overwrite the previous commands and the schedule workflow, correspondingly.

Algorithm 1: Algorithm for order creation.

Inputs: set of $l \in N$ devices (D^i), public keys (D_{PK}^i)
Variables: \mathcal{P} a puzzle, z a command.

```

1 Function create( $c_l, D^N$ )
2 begin
3   for  $\forall(i, j) \exists(D^i, D^j) \in H$  do
4     schedule(( $D^i, D^{i+1}$ ), ( $D^j, D^{j+1}$ ))
5     for  $\forall(i, i+1)$  do
6        $z = (s_{on} \vee s_{off}) \wedge (\hat{t}_i \leq \hat{t}_{i+1})$ 
7       generate( $\mathcal{P}^i$ ) = ( $n, a, \hat{t}_i, E_z, E_k$ )
8        $m^i = \text{enc}(\mathcal{P}^i, D_{PK}^i)$ 
9     end for
10    end for
11  return  $c_l = ((m^i, m^{i+1}), (m^j, m^{j+1}))$  end

```

- *Public parameter generator.* $\text{param}(1^\lambda) \rightarrow (n, a)$. $\text{param}(1^\lambda)$ initializes the prime integer factors $n = pq$ (where p and q are two large prime numbers) and random value a for the puzzle creation.
- *Puzzle generator.* $\text{puzgen}(n, a, \hat{t}, E_z, E_k) \rightarrow \mathcal{P}$. $\text{puzgen}(n, a, \hat{t}, E_z, E_k)$ selects the input values as target time for commands execution \hat{t} , encrypted command $E_z(z, k)$, encrypted key $E_k(a, \hat{t}, n, k)$ and outputs a puzzle \mathcal{P} for each device.
- *Follower* $D^i(\mathcal{P}^i, SK) \rightarrow z(\hat{t}_i, t_{com}^i)$ completes the puzzle \mathcal{P}^i and executes the command within the half-open interval, i.e., no earlier than \hat{t}_i but earlier than or at t_{com}^i .
- *Owner* $O(\mathcal{P}, \phi(n), PK^i) \rightarrow \text{accept}(t_{com}^i \geq \hat{t}_i)$ accepts the timely command execution at each device using $\phi(n)$.

Setup and key distribution. The manufacturing authority initializes a unique identity for the owner, hub, and, home devices by using the secure identity distribution function, say $\text{Init}(1^\lambda) \rightarrow \text{identity}$, where 1^λ is the security parameter that generates a unique *identity* for each entity. The certificate authority verifies that each of these devices knows the private key paired to the public key proposed for certification as: the homeowner (O_{SK}, O_{PK}), the hub (H_{SK}, H_{PK}), and i^{th} home device (D_{SK}^i, D_{PK}^i), possesses a valid key pair.

Step 1: Order creation: schedule creation at the owner. The homeowner O first creates a schedule, say *Schedule*, for device actuation. The creation of schedule is inherently specific to the preferences of the owner on a day to day basis and can include all or a subset of the home devices. The schedule creation does not require interaction with any device D^i or the hub H . Below we show a partially ordered timeline/schedule of four devices:

$$\text{Schedule} = ((D^1, D^2), (D^3, D^4))$$

Where only the elements of the same subset are comparable, e.g., D^1 with D^2 , and, D^3 with D^4 , based on the timeline. The owner converts this schedule to a verifiable device ordering (see Algorithm 1), before sending it to hub. The function $\text{create}(c_l, D^N)$ of Algorithm 1 converts a schedule for devices into the partially ordered sets of commands of length c_l , where $l \subseteq N$. Line 3 considers a pair⁹ of devices in the topology of the hub. Line 4 creates a mutually dependent schedule for the pair of devices with temporal dependency. Lines 5 and 6 consider all of these mutually dependent

⁹We consider only pairs of devices, to simply demonstrate the relative ordering. However, a different subset may contain as large as the total number of devices.

Algorithm 2: Algorithm for chaining.

Inputs: set of devices (D^N), user-defined schedule
 ($(D^i, D^{i+1}), (D^j, D^{j+1})$)

```

1 Function chain( $D^i, D^j, \mathcal{R}$ ) begin
2   for  $\forall(i, j, \mathcal{R}) \in \text{schedule}(D^i, D^j)$ ;  $i\mathcal{R}j = i < j$  do
3      ${}^n P_n(D^i)$ 
4     for  $\forall\{D^i\}_n! \wedge i\mathcal{R}j = \text{true}$  do
5       return  ${}^n P'_n\{D^i\}_{i=1}^n$ 
6     end for
7   end

```

pairs, decide the state of command as $z = (s_{on} \vee s_{off})$, and then, generate a unique puzzle \mathcal{P}^i for each device in Line 7. Here, the puzzle message contains a tuple of variables ($n, a, \hat{t}_i, E_z, E_k$), where n is the product of two large prime numbers p and q , a is a random number, \hat{t}_i is the time-complexity of the puzzle, E_z is the encrypted command z using key k , and E_k is the encrypted key k . Line 8 encrypts each puzzle \mathcal{P}^i into a message m^i for a device D^i using the public key of the device D_{PK}^i . Line 9 returns an assembled order c_l that contains an encrypted message for each device corresponding to the mutually dependent devices in the schedule. This ends the creation of a relative order for the chosen set of devices. Next, the homeowner sends securely this order, say *Order*, to the hub, as follows:

$$\text{Order} = (O_{id}, H_{id}, c_l, \text{Sign}(\mathcal{H}, O_{SK}))$$

i.e., the homeowner sends its identity (O_{id}), the identity of the hub (H_{id}), and encrypted order of commands (c_l) along with hash digest of all three attributes.

Step 2: Token generation by the hub and token delivery to devices. On receiving the partial order of commands from the owner, the hub verifies the sender by computing a local hash digest \mathcal{H}' over (O_{id}, H_{id}, c_l). Also the hub verifies the signature using O_{PK} and compares the received hash digest \mathcal{H} with the locally computed hash digest \mathcal{H}' . If $\mathcal{H} = \mathcal{H}'$ then the hub accepts this order.

After the verification of order origination, the hub creates a token that is used for order delivery (in this step) and for data collection generated by devices (step 4). The token has three fields: command field c_l , data field, and toggle bit string b_{toggle} . Every token field has sensitive information regarding the device activity. Therefore, we assume that the token is encrypted using a shared symmetric key k_s among the devices and the controller hub.

$$\mathcal{T} = E((c_l || \text{Data field} || b_{toggle}), k_s)$$

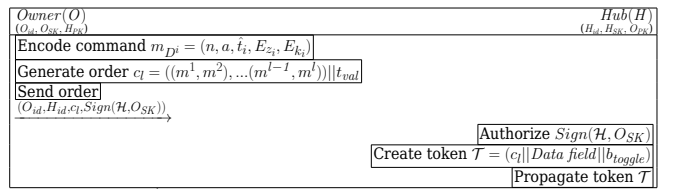


Figure 3: Order creation.

The hub, then, forwards the token among all devices in the topology even if a device was not included in the schedule, as shown in Figure 3. On receiving the token, each device decodes the corresponding command in token and forwards the token to the next

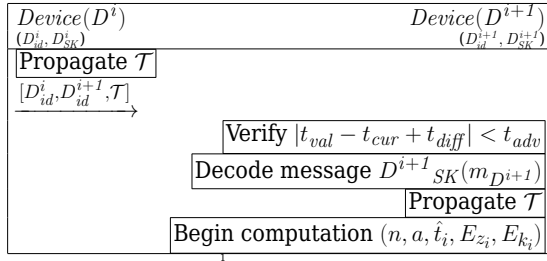


Figure 4: Token circulation.

peer device in the topology. The next peer device is chosen as per the topology underneath. Recall that the devices are connected in a pre-defined topology¹⁰ such as in an unidirectional ring, bidirectional ring, star (fault-tolerant), grid, mesh or hybrid setting. Note that the token rotates constantly in the topology (see Figure 4).

Algorithm 2 explains the linear ordering of devices in function $chain(D^i, D^j, \mathcal{R})$. The linear ordering condition requires that each device D^i and D^j must satisfy: the exact same mutual ordering or relation $\mathcal{R} = \{\leq\}$ as in c_l . The line 2 selects each pair (i, j) of device that is paired under relation $\mathcal{R} = \{\leq\}$ in the original schedule. Line 3 enumerates all possible permutations of these devices, say ${}^n P_n(D^i)$ where $i \in N$. Line 4 selects one permuted order ${}^n P'_n$ of devices (from the total number of permutations) such that the precedence relation still holds true, i.e., $i\mathcal{R}j = true$. Finally, in line 5, the selected topological order ${}^n P'_n$ is returned from all possible permuted orders ${}^n P_n(D^i)$ of the devices.

Step 3: Order retrieval and puzzle computation at the devices.

Note that the activation sequence of devices is released ahead of their actual activation; however, we need to restrict them not to execute the command before prescribed time. Thus, in this step, devices perform a pre-defined computation task (as detailed in Figure 5) to unlock and execute the owner-defined command.

It must be noted that unlocking the command is as necessary as unlocking the command within the prescribed duration, i.e., knowing the particular order of a device in the overall sequence. A device receives the computation task in the form of a puzzle as soon as the order is delivered in step 2. Subsequently, the device begins the computation task if the puzzle validity period has not expired yet. The time-bound during which a device is restricted to begin, as well as, end the computation task cannot be compressed unless the device possesses a distinguisher for the factoring problem. Therefore, the verification that the secure computation task is crucial. This can be verified through the Euler's totient as a trapdoor for factoring n inside the puzzle. Note that the value of n is public and the value of $\phi(n)$ is kept secret. Therefore, computing $\phi(n)$ from n is as hard as integer factorization. In addition, without the knowledge of $\phi(n)$ the computation time for E_k is directly dependent on \hat{t} time-consuming square operations.

Algorithm 3 explains the verification of time-bounded commands in function $verify(D^i, \hat{t}_i)$. Line 3, considers all devices that are part of the current schedule. In order to verify that a specific device has executed the command within the pre-defined interval, O securely

¹⁰It must be noted that from the practical deployment aspect it is difficult to connect these smart devices in a ring topology unless the devices belong to the same OEMs, e.g., Apple HomeKit. Therefore, a star or a grid topology can be used to combat the single point of failure and device heterogeneity in the current scenario.

The puzzle messages:

$$\begin{aligned} n &= pq \\ \hat{t} &= St' \\ E_z &= enc(z, k) \\ E_k &= k + a^{2^i} \pmod n \end{aligned}$$

The puzzle computation (\mathcal{P}):

- (1) Initially, \mathcal{P} receives as input a secret key k and encrypt the original message z denoted as E_z . It must be noticed that each individual command corresponding to a device D^i is secured in the form of an encrypted message m_i . Furthermore, a cascaded command as a whole contains multiple messages of these types.
- (2) Subsequently, \mathcal{P} receives inputs as a secret key k , random number a , puzzle difficulty level \hat{t} , modulus n , and then generate E_k . The puzzle computation relies on the secrecy of key k used to encrypt a secret message $E_z = (z, k)$. Also, $\hat{t} = S * t'$ denotes the difficulty level of the puzzle for a specific device which can perform S number of square operations per second and t' is the time to decrypt the message using a regular encryption scheme.
- (3) The puzzle \mathcal{P} includes the tuple $(n, a, \hat{t}, E_z, E_k)$ for which the computation task is to be solved. Evidently, the recipient of this puzzle would have to spend at least \hat{t} amount of time to complete the computation task and reveal the key k .

Figure 5: Puzzle computation.

Algorithm 3: Algorithm for delay verification.

Inputs: a puzzle (\mathcal{P}) with a set of public variables $(n, a, \hat{t}_i, z_i, k_i)$

```

1 Function  $verify(D^i, \hat{t}_i)$ 
2 begin
3   for  $\forall(i, j) \in schedule((D^i, D^{i+1}), (D^j, D^{j+1}))$  do
4     if  $a^{2^i} \pmod n \equiv a^{2^i} \pmod{\phi(n)} \pmod n$  then
5       if  $(t_{com}^i \leq t_{com}^{i+1}) \wedge (t_{com}^j \leq t_{com}^{j+1})$  then
6         return True
7       end if
8     end for
9   return False end

```

pre-computes the Euler's totient $\phi(n) = (p-1)(q-1)$ such that p, q is discarded after computing the n and $\phi(n)$. Line 4, verifies the time bound for each of these devices, such that

$$a^{2^i} \pmod n \equiv a^{2^i} \pmod{\phi(n)} \pmod n$$

Line 5 compares the execution order of devices that have happened as a result of puzzle computation and returns true in line 6 if it is a total order.

Step 4: Data generation at the devices. Once the devices have completed the puzzle computation, they generate the data as a result of the command execution. These home devices are bound to upload the locally generated data to the hub. In our scheme, the token contains an anonymous data field to securely transmit the device generated data to the hub. We use bitwise (b) XOR padding to overwrite the random data in the token data field as:

Overwrite data (b_o) = *Random data* (b_r) \oplus *Generated data* (b_g)

It is cryptographically hard to distinguish the presence of random data from the device generated data as stored inside the token. Note that our threat model does not consider the ISP or DNS level threats, therefore, devices are only assumed to securely generate and anonymously dispatch the data to the hub and combat any passive learning attacks within the physical periphery of the home. *Collision*: Our token circulation strategy and the token structure are primarily for smart home scenarios, where we assume that the single token field can accommodate the peak hour traffic. However, in case the peak hour traffic exceeds and multiple devices request for data upload, for example in a multi-tenant building, then to avoid the collision situation more data fields are required. A simple approach is to create sub-fields inside the data field such that each sub-field belongs to a unique device. Thus, each device can fairly utilize the data upload capacity in any round during the token circulation.

3.3 Time Analysis

The time spent during the token circulation and puzzle computation is directly proportional to the number of devices connected in the network. For example, token begins at time (t_{beg}^H) at the hub and completes the first round of token circulation at time (t_{end}^H). The time spent at i th device is ($t_{fwd}^i - t_{rcv}^i$) that receive the token at time (t_{rcv}^i) and forward it to next device at time (t_{fwd}^i). Therefore, the total time spent in one round of token circulation:

$$t_{sum} = (t_{end}^H - t_{beg}^H) - (t_{fwd}^i - t_{rcv}^i)_{i=1}^N$$

Note that the token circulation time is sequenced and linear w.r.t. the number of devices. While the puzzle computation time t_{com}^i varies independently among all devices. So the puzzle computation time at i th device is $t_{com}^i \approx \hat{t}^i$. In order to optimize the puzzle computation time and still retain the verifiable guarantees on the artificial delay, we consider three type of puzzles.

- *For comparable devices*: Each pair of comparable devices in the topology requires that \hat{t} 's are at least $(N-1)(t_{fwd}^N - t_{fwd}^{N-1})$ apart. The devices forward the token before beginning the local computation task. Any two adjacent devices (D^{N-1}, D^N) that begin the computation after forwarding the token, must possess:

$$|\hat{t}^N - \hat{t}^{N-1}| \geq (N-1)(t_{fwd}^N - t_{fwd}^{N-1})$$

- *For incomparable devices*: The set of incomparable devices require that \hat{t} 's are exactly $(j-i)(t_{fwd}^j - t_{fwd}^i)$ apart. Every time a device D^i forwards a token to D^{i+1} it jumps $(t_{fwd}^{i+1} - t_{fwd}^i)$ ahead on the computation timeline with respect to next device due to token propagation delay. Therefore, in order to provide an identical time of actuation for all incomparable devices:

$$|\hat{t}^i - \hat{t}^j| = (j-i)(t_{fwd}^j - t_{fwd}^i)$$

The total number of slots required is $(N-k)+1$ where k represents the number of comparable devices. In particular, each comparable device requires a unique and non-overlapping $|\hat{t}|$ w.r.t. other comparable devices; while each incomparable device can be scheduled for an identical and overlapping $|\hat{t}|$.

Token frequency: The token frequency is a crucial attribute from the perspective of how early a user can decide the schedule for all N devices and, how many data upload requests are received during the peak hours. The frequency of token circulation can be either *fixed* or *random*. Let us assume a *fixed* slot i between any two consecutive rounds of the token circulation such that the token begins a new round at every i th unit of time. The optimal length of the slot is the same as the maximum \hat{t}^i in any schedule.

$$\text{Slot length} = \max\{\hat{t}^i\}_{i=1}^{l \in N}$$

For example, if $(\hat{t}^1, \hat{t}^2, \dots, \hat{t}^{l-1}, \hat{t}^l)$ is the time-bound for l devices in any scheduled workflow then the slot length is same as the farthest possible device on the timeline of a scheduled workflow.

Table 3 represents the cost comparison based on mathematical operations such as encryption (E), decryption (D), signature generation and verification (S), exclusive-OR (XOR), hashing (H), scalar multiplication (M), and modular exponentiation (Me). The scheme in [6] imposes a relative overhead such that (p_m) number of masking packets are required per traffic flow in case the traffic flow is lesser than a pre-defined threshold value. Therefore, the total overhead per traffic flow is (Cp_m) where C is the communication overhead per packet. Similarly, if the traffic flow is above the threshold value then those excess packets p'_m are delayed and stored inside a queue. Therefore, the total latency per traffic flow is (Tp'_m) where T is the latency overhead per packet. As shown here, that our proposed scheme requires the minimum number of operations. Further, in our approach the computational complexity at devices is variable and it depends on the required number of modular exponentiations, e.g., \hat{t} , as initialized by the owner.

4 SECURITY ANALYSIS

This section provides the security analysis for proposed scheme. We first model the security experiment, below, like the standard security model.

ATTACK GAME 1. Let I be the order-preserving protocol between the challenger and adversary \mathcal{A} then the attack game works as:

- *Public parameter generation*: The challenger generates (n, a) using $param(1^\lambda)$.
- *Puzzle generation*: The challenger generates \mathcal{P} using $puzgen(n, a, \hat{t}, E_k, E_z)$.
- *Query phase*: An adversary attempts to attack I through token query given the access to a recently generated token \mathcal{T}' . The adversary sends a value \hat{t} to the challenger. The challenger generates the corresponding puzzle \mathcal{P} and adds in c_I . The follower devices receive c_I , extract the unique puzzle, and execute the command. The challenger then sends \mathcal{T}' to the adversary.
- *State identification attempt*: The adversary attempts to retrieve the intermediate state of the computation task and attempts to solve the puzzle earlier than device through $t_{com}^{\mathcal{A}}$ for the same \mathcal{P} in \mathcal{T}' , such that $t_{com}^{\mathcal{A}}$ is lower than the original \hat{t} . An adversary \mathcal{A} wins the game, if $t_{com}^{\mathcal{A}} < \hat{t}_i$ and the owner outputs *accept*. The probabilistic advantage of the adversary, $Adv(\mathcal{A})$, for winning the game is:

$$Adv(\mathcal{A}) = Pr[t_{com}^{\mathcal{A}} < \hat{t}_i]$$

| Protocol | Cost at device | Cost at hub | Cost at owner |
|-----------------|---------------------------|-------------------------|-----------------------------------|
| Proposed scheme | $(2D + 1XOR + \hat{t}Me)$ | $(1S + 1E + 1XOR)$ | $(1S + 1E + \hat{t}mod\phi(n)Me)$ |
| Scheme [11] | $(3H + 7XOR + 1E + 1D)$ | $(5H + 8XOR + 1E + 1D)$ | - |
| Scheme [13] | $(3M + 2H + 4XOR)$ | $(1H + 4XOR)$ | - |
| Scheme [6] | - | $(Cp_m + Tp'_m)$ | - |

Table 3: Cost comparison between our scheme and existing schemes [6, 11, 13].

We present a sequence of games as Game 0 to Game 2. Each Game i shows that the advantage of an adversary $Pr[t_{com}^{\mathcal{A}} < \hat{t}_i]$ is negligibly small. Similarly, each subsequent game Game $(i+1)$ is produced through previous game such that the changes in secret parameters remain indistinguishable to the adversary. Therefore, the advantage of an adversary through changes in secret parameters (i.e., transitioning from one game to another) remains negligibly small. If $Pr[\mathcal{A}(i) \rightarrow 0] - Pr[\mathcal{A}(i+1) \rightarrow 0]$ is non-negligible then that adversary can be used as a distinguisher or as a solver for the integer factorization problem in our scheme; where $Pr[\mathcal{A}(i)]$ and $Pr[\mathcal{A}(i+1)]$ represent the probability of adversary winning the Game i and Game $(i+1)$, respectively. The Game 0 represents the original attack such that $(t_{com}^{\mathcal{A}} = \hat{t}_i)$ and the artificial delay before the command execution is kept null. The Game 1 represents the attack with enhanced t' while $(t_{com}^{\mathcal{A}} < t' \wedge \hat{t}_i)$. Similarly, the Game 2 represents the attack with general t' while $(t_{com}^{\mathcal{A}} = t')$ but $(t_{com}^{\mathcal{A}} < \hat{t}_i)$.

Game 0: **[Record attack]** Let us assume that the puzzle P_1 contains $\hat{t} = 0$ then the device must perform only one iteration to compute and decode the enciphered command. However, an adversary cannot distinguish an early puzzle such as P_1 from a delayed puzzle such as P_2 for which $\hat{t} > 0$. In the token query phase, an adversary gathers the token transcripts for a known value of \hat{t} .

Experiment $\text{Exp}_{\mathcal{A}}^{t_{com}^{\mathcal{A}}=\hat{t}_i}$
 let $c_l((m^1, m^2), \dots, (m^{l-1}, m^l)) \leftarrow \mathcal{A}(\mathcal{T})$
 generate $m^i(\hat{t})$ at random
 $(m^k(\hat{t})) \leftarrow \mathcal{A}(c_l((m^1, m^2), \dots, (m^{l-1}, m^l)))$
 if $(m^k(\hat{t}) = m^i(\hat{t}))$
 return 1
 else return 0

It is computationally hard to distinguish between encrypted commands and to identify the command that carries known \hat{t} within time $t_{com}^{\mathcal{A}} = \hat{t}_i$. An adversary can distinguish the commands with the advantage

$$Adv(\mathcal{A})_{\text{Game } 0} = Pr[m^i(\hat{t}) \leftarrow c_l((m^1, m^2), \dots, (m^{l-1}, m^l))] \leq \epsilon$$

Game 1: **[Clone attack with lesser t']** Let us assume that the adversary computes the puzzle in time $t_{com}^{\mathcal{A}} < t' \wedge \hat{t}_i$ where the unit time capacity t' of the adversary is slower. Therefore, the advantage of the adversary depends on the probability to compute the \hat{t} square operations faster than the home device. This requires that the adversary can solve the prime factors for modulus n .

$$Adv(\mathcal{A})_{\text{Game } 1} = Pr[(p, q) \leftarrow (n)] \leq \epsilon$$

Game 2: **[Clone attack with general t']** Let us assume that the adversary can compute as fast as the home device, i.e., t' . In particular, the adversary can also perform S number of square operations per second. The probability $Pr[t' = t_{com}^{\mathcal{A}} < \hat{t}_i]$ that an adversary \mathcal{A} can solve a puzzle with the difficulty level \hat{t} in lesser time than

| Properties | [11] | [13] | [6] | Our scheme |
|--------------------------|------|------|-----|------------|
| Upstream direction | ✓ | ✓ | ✓ | ✓ |
| Downstream direction | ✓ | × | ✓ | ✓ |
| Verifiable delay | × | × | × | ✓ |
| Partial ordering | × | × | × | ✓ |
| Total ordering | × | × | × | ✓ |
| Privacy | × | × | ✓ | ✓ |
| Passive attack resistant | × | × | ✓ | ✓ |
| Active attack resistant | ✓ | ✓ | × | ✓ |

Table 4: Comparison between our and existing schemes.

the home device is negligibly small. Since the adversary must compute S' number of operations for each S operations at home device where $(S' - S > \epsilon)$:

$$Adv(\mathcal{A})_{\text{Game } 2} = Pr[t_{com}^{\mathcal{A}}(S't') \leftarrow \hat{t}_i(S't')] \leq \epsilon$$

In this sequence of games Game 0 through Game 2 the total advantage of the adversary $Adv(\mathcal{A})$ depends on the sum of the probability to win each of these games.

$$Pr[t_{com}^{\mathcal{A}} = \hat{t}_i] + Pr[t_{com}^{\mathcal{A}} < t' \wedge \hat{t}_i] + Pr[t' = t_{com}^{\mathcal{A}} < \hat{t}_i] \leq \epsilon$$

Overall, the advantage of the adversary is proportional to the availability of computational resources to solve the puzzle for all devices in parallel. Similarly, the advantage of the adversary with respect to a single puzzle and a single device are proportional to the availability of computational resources to solve the inherently sequential operations of that individual puzzle. Therefore, the total advantage of an adversary to clone the entire timeline depends on the total computational power for both, the parallel and the sequential operations to decrypt the command ahead of time.

In Table 4 a comparison is shown between our proposed scheme and the existing work. The comparison is based on the data flow direction, ordering, verification of ordering, privacy violation, and attack resistance.

5 EXPERIMENTAL EVALUATION

This section evaluates our proposed system using our prototype implementation. We describe the mock-up testing IoT application, experimental setup and overall results from our experiments.

5.1 Experiment Setup

In order to demonstrate the effectiveness and performance of our proposed architecture, we developed the prototype implementation and setup the testbed in our lab (as shown in Figure 6). This proof-of-concept prototype implements the protocols described in Section 3 and a test application with Python. In this mock-up IoT application, a device awaits and executes two types of command given by the homeowner. The “Set” command will change a variable in the program of the target device while the “Read” command will

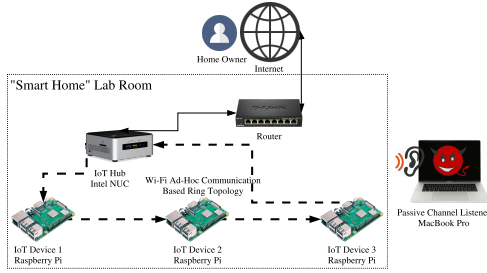


Figure 6: Experimental testbed in our lab

require the device to send the variable together with device status, e.g., RAM and CPU usages, back to the homeowner. This mock-up application is created to simulate two-way communication between the homeowner and devices. Further, note that the application runs on top of the approach, we proposed in this paper.

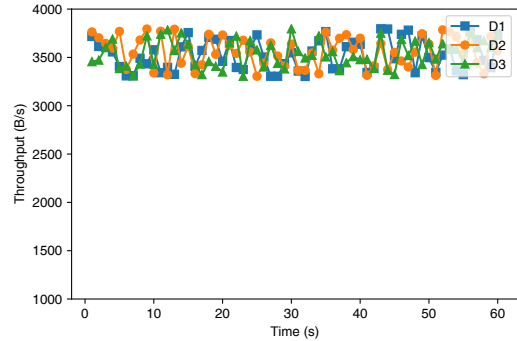
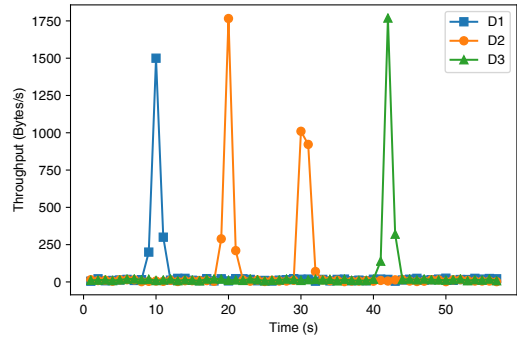
Figure 6 depicts the architecture and configuration of the smart home testbed. An Intel NUC system is programmed to work as the hub H that forwards information between the homeowner and smart devices. The hub H equips with two network interfaces: (i) Ethernet interface that has connections to receive the homeowner-defined schedule and sends the data of smart devices to the homeowner, and (ii) Wi-Fi interface that is used to communicate with smart home devices. Smart home device programs are deployed on three Raspberry Pis (D^1, D^2, D^3) (3rd Gen B+ Model) that are equipped with built-in Wi-Fi interface. Wi-Fi interfaces on the hub H and the devices (D^1, D^2, D^3) are configured to work in Wi-Fi ad-hoc mode [4], which enables direct device-to-device communication. All Wi-Fi interfaces are configured with pre-defined Wi-Fi channel, static IP address, and routing information to have a ring topology. A MacBook Pro with its Wi-Fi interface is deployed in a different room next to the testbed performs as an adversary, who listens and dumps all channel activities on the pre-defined channel into the pcap (packet capture) file for future analysis.

5.2 Results

Based on the testbed described above, we performed different experiments to evaluate the proposed system and the approach. We first validate our system to check whether it could prevent the adversary from learning device activity from channel activity or not. Then, we explore the performance of the proposed ring topology communication with a set of experiments.

Decoupling channel activities from device activities. To evaluate the effectiveness that our system protects against the passive channel listeners, we first defined a sequence of sample user commands, e.g., D^1 : Set, D^2 : Set, D^2 : Read, and D^3 : Read. In a one-minute experiment, these commands will be issued in 10 seconds intervals. We performed the experiment by executing the above-mentioned sequence in our proposed system and also over Wi-Fi infrastructure network without a ring topology to compare with. The channel activities are recorded by the passive channel listener laptop deployed near our testbed.

Figure 7a shows the passive adversary’s view, i.e., which device receives a message from the homeowner at which time due to channel activities in the experiment. It is clear that each time the user sends a command to devices or devices send data to the user, there will be a peak in the channel activity. Thus, the adversary infers



(b) Devices working with our proposed system

Figure 7: The adversarial view due to observing channel activities

the device activity and user-device interaction from channel activity. In contrast, Figure 7b shows the effectiveness of our proposed approach. Note that the channel activity patterns are completely eliminated, due to the token ring communication.

Communication latency. Instead of sending individual commands or data to/from devices or hub, the commands and data in our system are encapsulated in tokens and transmitted in a ring topology, which will incur additional communication latency. We performed experiments to evaluate the impact of increasing latency as the number of devices in the ring topology increases. Since we only have a very limited number of *real* devices, we modified our protocol to simulate the scenario that includes a large number of devices to investigate an impact on communication latency. To achieve this goal, we add a counter in each token. When the hub generates the token, it sets the counter equals to the number of devices we want to simulate in the experiments. This counter is decreased by one when each device receives and forwards the token to the next device. When the last device (D^3 in our testbed) in the ring topology receives the token, it checks the value of the current counter. If the value of the counter is more than zero, D^3 forwards the token to the first device (D^1) to extend the ring topology. If the counter number is less than or equal to zero, the last device forwards the token back to the hub to complete a single round of the token. Here, since each device will receive the same token multiple times, we also need to prevent the device from solving same puzzles and executing same commands multiple times. To do so, a unique token ID is added to each token. Thus, the device solves the puzzle and executes the command only when the device gets a new token ID.

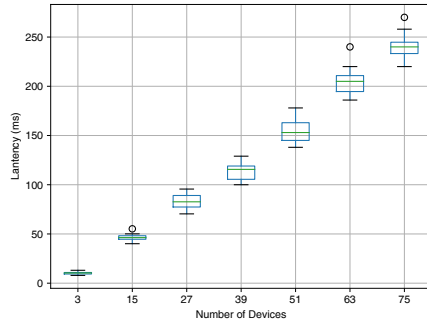


Figure 8: Impact of ring topology on the latency

Figure 8 shows that both mean and variation of the latency increase as more devices added into the ring topology. The mean of latency rises linearly at the beginning as each additional hop in the ring topology introduces more latency. After around 39 devices in the ring topology, the mean latency starts rising faster, since the length of token also increases with the growth of a number of devices. Consequently, it may take more time at devices to transmit the token to the next hop. Of course, the ring latency is not affected, when there is a few devices, since the number of commands and toggle bit strings decrease in the token as an decreasing number of devices. The variation of latency gets larger since it is more likely that the token transmission at more hops get delayed or re-transmitted because of unexpected interference or system lag at the devices.

| # Devices | 3 | 27 | 51 | 63 | 75 |
|-----------------------|------|------|------|------|-------|
| Avg. Token Len(bytes) | 1807 | 5024 | 8189 | 9757 | 11305 |

Table 5: Average token length for ring topology size

Communication and computation overheads. It is also important to understand the additional overheads incurred due to our approach. We first measure how the length of token increases as the number of devices grows. Table 5 shows that the average token length grows linearly. We also use a USB power meter to measure the power consumption of the Raspberry Pi in different working states. Table 6 shows that our system introduces 63% more power consumption to completely eliminate the channel activity patterns of all devices.

6 RELATED WORK

There exist various IoT frameworks such as *Apple HomeKit*, *Smart-Thing*, *Azure IoT Suite*, *IBM Watson IoT* platform, *Brillo/Weave* platform by Google, *Calvin IoT* platform by Ericsson, *ARM mbed IoT* platform, *Kura IoT* project by Eclipse, interested readers may refer [3, 7] for more details. Overall, the smart home communication [8, 10] is a network component that executes commands based on contextual factors. To the best of our knowledge, none of the work [5, 6] highlight the significance of secure device ordering in a smart home scenario. We highlight the presence of a pattern among smart home devices such that a partial ordering on device activity is observed on daily basis. The authors in [6] have presented a privacy-preserving traffic shaping scheme to mask the channel activity and thereby the device or user activity at the ISP level. According to the scheme, if the shaped traffic rate is lower than the device traffic then the packets are queued, and if the shaped traffic rate is higher than the device traffic then the dummy packets are added to

| State | Idle | IoT App w/o Ring Sys. | IoT App w/ Ring Sys. |
|------------|-------|-----------------------|----------------------|
| Avg. Power | 2.25W | 3.03W | 4.91W |

Table 6: Average energy consumption of devices in different working states

cover the original traffic rate variations. However, these techniques do not avoid the inferences on device activity pattern due to the straightforward binding between the channel activity and device activity. Our scheme decouples the channel activity from the device activity such that a communication activity over the channel at any given time cannot be coupled with a specific device activity or the user activity. The recent work [11, 13] provides a security framework for home devices and guarantees message anonymity and unlinkability during the communication sessions from hub to the device. The scheme is based on authentication and one-time session key agreement in a 3-way handshake protocol. However, as mentioned, authentication and encrypted message cannot prevent the inferences over the communication activity.

7 CONCLUSION

This paper focuses on the security and privacy challenges in smart homes that facilitate the execution of multiple workflows. In particular, our solution avoids any inference attacks regarding these workflows that can reveal the user activity pattern, both in the past and the future. The primary source of these inference attacks is the ability to learn communication patterns among the smart home devices. The channel activity and the corresponding device activities are also sensitive from the user’s privacy perspective. Therefore, decoupling the channel activity from the device activity is essential to hide the execution of scheduled workflows.

REFERENCES

- [1] In *Insight’ into Home Automation Reveals Vulnerability in Simple IoT Product*, available at URL: <https://securingtomorrow.mcafee.com/mcafee-labs/insight-into-home-automation-reveals-vulnerability-in-simple-iot-product>.
- [2] In *IEEE OUI (Organizationally Unique Identifier)*, available at URL: <http://standards-oui.ieee.org/oui.txt>.
- [3] M. Ammar et al. Internet of things: A survey on the security of IoT frameworks. *JISA*, 38:8 – 27, 2018.
- [4] G. Anastasi et al. Wi-fi in ad hoc mode: a measurement study. In *PerCom*, pages 145–154, 2004.
- [5] N. Apthorpe et al. Closing the blinds: Four strategies for protecting smart home privacy from network observers. *CoRR*, abs/1705.06809, 2017.
- [6] N. Apthorpe et al. Spying on the smart home: Privacy attacks and defenses on encrypted IoT traffic. *CoRR*, abs/1708.05044, 2017.
- [7] T. T. Doan et al. Towards a resilient smart home. In *IoT S&P*, pages 15–21, 2018.
- [8] S. M. D’Souza et al. Time-based coordination in geo-distributed cyber-physical systems. In *HotCloud*, 2017.
- [9] S. Feng et al. Smart home: Cognitive interactive people-centric internet of things. *IEEE Communications Magazine*, 55(2):34–39, 2017.
- [10] D. N. Kalofonos et al. Intuisec: A framework for intuitive user interaction with smart home security using mobile devices. In *PIMRC*, pages 1–5, 2007.
- [11] P. Kumar et al. Anonymous secure framework in connected smart home environments. *IEEE Trans. Information Forensics and Security*, 12(4):968–979, 2017.
- [12] R. L. Rivest et al. Time-lock puzzles and timed-release crypto. Technical report, MIT/LCS/TR-684, MIT Lab for Computer Science, 1996.
- [13] J. Shen et al. Secure data uploading scheme for a smart home system. *Information Sciences*, 453:186 – 197, 2018.
- [14] G. J. Watson et al. Lost: Location based storage. In *ACM Workshop on Cloud Computing Security*, pages 59–70, 2012.